

What Is Claimed Is:

1 1. A method for bypassing use of a protocol checksum during
2 communications across a reliable network link, comprising:
3 configuring a communication system to bypass use of the checksum during
4 communications across the reliable network link;
5 receiving an outbound packet to be transmitted to a destination across the
6 reliable network link; and
7 sending the outbound packet to the destination across the reliable network
8 link without computing the checksum for the outbound packet.

1 2. The method of claim 1, wherein configuring the communication
2 system to bypass the checksum involves informing a protocol stack within the
3 communication system that network interface hardware for the communication
4 system is capable of computing the checksum, so that the protocol stack does not
5 compute the checksum.

1 3. The method of claim 1, further comprising:
2 determining whether the outbound packet is directed to a valid destination
3 that is eligible for checksum bypassing;
4 if the outbound packet is not directed to a valid destination,
5 computing the checksum for the outbound packet, and
6 inserting the checksum into the outbound packet.

1 4. The method of claim 3, wherein the checksum is computed by a
2 driver associated with network interface hardware for the communication system.

1 5. The method of claim 1, further comprising:
2 receiving an inbound packet from a source across the reliable network
3 link; and
4 accepting the inbound packet without re-computing the checksum;
5 wherein re-computation of the checksum is required by the communication
6 protocol to verify that the inbound packet was received without errors.

1 6. The method of claim 5, wherein accepting the inbound packet
2 without re-computing the checksum involves:
3 communicating a default checksum value to a protocol stack within the
4 communication system;
5 wherein the default checksum value matches the default checksum value
6 contained within a checksum field of the inbound packet;
7 whereby the protocol stack will match the default checksum value with the
8 checksum field of the inbound packet and will consequently accept the inbound
9 packet.

1 7. The method of claim 6, wherein accepting the inbound packet
2 without re-computing the checksum additionally involves inserting the default
3 checksum value into the checksum field of the inbound packet.

1 8. The method of claim 1, wherein the communication protocol
2 includes one of:
3 Transmission Protocol (TCP);
4 Internet Protocol (IP); and
5 User Datagram Protocol (UDP).

1 9. The method of claim 1, wherein the reliable network link adheres
2 to the InifiBand standard.

1 10. The method of claim 2,
2 wherein the checksum is a TCP checksum; and
3 wherein the protocol stack is an IP stack.

1 11. A computer-readable storage medium storing instructions that
2 when executed by a computer cause the computer to perform a method for
3 bypassing use of a protocol checksum during communications across a reliable
4 network link, the method comprising:

5 configuring a communication system to bypass use of the checksum during
6 communications across the reliable network link;
7 receiving an outbound packet to be transmitted to a destination across the
8 reliable network link; and
9 sending the outbound packet to the destination across the reliable network
10 link without computing the checksum for the outbound packet.

1 12. The computer-readable storage medium of claim 11, wherein
2 configuring the communication system to bypass the checksum involves
3 informing a protocol stack within the communication system that network
4 interface hardware for the communication system is capable of computing the
5 checksum, so that the protocol stack does not compute the checksum.

1 13. The computer-readable storage medium of claim 11, wherein the
2 method further comprises:

3 determining whether the outbound packet is directed to a valid destination
4 that is eligible for checksum bypassing;
5 if the outbound packet is not directed to a valid destination,
6 computing the checksum for the outbound packet, and
7 inserting the checksum into the outbound packet.

1 14. The computer-readable storage medium of claim 13, wherein the
2 checksum is computed by a driver associated with network interface hardware for
3 the communication system.

1 15. The computer-readable storage medium of claim 11, wherein the
2 method further comprises:
3 receiving an inbound packet from a source across the reliable network
4 link; and
5 accepting the inbound packet without re-computing the checksum;
6 wherein re-computation of the checksum is required by the communication
7 protocol to verify that the inbound packet was received without errors.

1 16. The computer-readable storage medium of claim 15, wherein
2 accepting the inbound packet without re-computing the checksum involves:
3 communicating a default checksum value to a protocol stack within the
4 communication system;
5 wherein the default checksum value matches the default checksum value
6 contained within a checksum field of the inbound packet;
7 whereby the protocol stack will match the default checksum value with the
8 checksum field of the inbound packet and will consequently accept the inbound
9 packet.

1 17. The computer-readable storage medium of claim 16, wherein
2 accepting the inbound packet without re-computing the checksum additionally
3 involves inserting the default checksum value into the checksum field of the
4 inbound packet.

1 18. The computer-readable storage medium of claim 11, wherein the
2 communication protocol includes one of:
3 Transmission Protocol (TCP);
4 Internet Protocol (IP); and
5 User Datagram Protocol (UDP).

1 19. The computer-readable storage medium of claim 11, wherein the
2 reliable network link adheres to the InifiBand standard.

1 20. The computer-readable storage medium of claim 12,
2 wherein the checksum is a TCP checksum; and
3 wherein the protocol stack is an IP stack.

1 21. An apparatus that bypasses use of a protocol checksum during
2 communications across a reliable network link, comprising:
3 a configuration mechanism that selectively configures a communication
4 system to bypass use of the checksum during communications across the reliable
5 network link;
6 a receiving mechanism that is configured to receive an outbound packet to
7 be transmitted to a destination across the reliable network link; and
8 a sending mechanism that is configured to send the outbound packet to the

9 destination across the reliable network link without computing the checksum for
10 the outbound packet.

1 22. The apparatus of claim 21, wherein the configuration mechanism
2 informs a protocol stack within the communication system that network interface
3 hardware for the communication system is capable of computing the checksum, so
4 that the protocol stack does not compute the checksum.

1 23. The apparatus of claim 21,
2 wherein the configuration mechanism is configured to determine whether
3 the outbound packet is directed to a valid destination that is eligible for checksum
4 bypassing; and
5 wherein if the outbound packet is not directed to a valid destination, the
6 configuration mechanism is configured to,
7 compute the checksum for the outbound packet, and to
8 insert the checksum into the outbound packet.

1 24. The apparatus of claim 23, wherein the checksum is computed by a
2 driver associated with network interface hardware for the communication system.

1 25. The apparatus of claim 21, wherein the receiving mechanism is
2 configured to:
3 receive an inbound packet from a source across the reliable network link;
4 and to
5 accept the inbound packet without re-computing the checksum;
6 wherein re-computation of the checksum is required by the communication
7 protocol to verify that the inbound packet was received without errors.

1 26. The apparatus of claim 25,
2 wherein the receiving mechanism is configured to communicate a default
3 checksum value to a protocol stack within the communication system; and
4 wherein the default checksum value matches the default checksum value
5 contained within a checksum field of the inbound packet;
6 whereby the protocol stack will match the default checksum value with the
7 checksum field of the inbound packet and will consequently accept the inbound
8 packet.

1 27. The apparatus of claim 26, wherein the receiving mechanism is
2 additionally configured to insert the default checksum value into the checksum
3 field of the inbound packet.

1 28. The apparatus of claim 21, wherein the communication protocol
2 includes one of:
3 Transmission Protocol (TCP);
4 Internet Protocol (IP); and
5 User Datagram Protocol (UDP).

1 29. The apparatus of claim 21, wherein the reliable network link
2 adheres to the InfiniBand standard.

1 30. The apparatus of claim 22,
2 wherein the checksum is a TCP checksum; and
3 wherein the protocol stack is an IP stack.